

An Historical Guide to the Ultracomputer Literature

Allan Gottlieb

Courant Institute, N.Y.U.
251 Mercer Street
New York, N.Y. 10012

Ultracomputer Note #36

October, 1981

When reading our reports, one should note that three distinct parallel processors are discussed: the "ultracomputer", a message passing machine; the "paracomputer", an idealized shared memory machine; and the "NYU Ultracomputer", a network-based shared memory machine that approximates a paracomputer. Moreover, since the NYU Ultracomputer has replaced the original ultracomputer as our preferred design, our later reports often "abbreviate" NYU Ultracomputer as ultracomputer. Finally, our basic coordination primitive has been called NEWVAL, replace-add, and fetch-and-add. We realize that such highly context sensitive language may be confusing and this guide is a (long overdue) attempt to help clarify the situation. In addition, by presenting a historical overview of the project I hope to group together those reports reflecting work done at roughly the same time (and hence using roughly the same terminology). Unfortunately, the report numbers and dates reflect only the order in which the work was written up in final form. Since the present note contains no new technical material, I hope that you will permit an informal exposition.

During the academic year 1978-1979, Jack Schwartz wrote the original paper on ultracomputers [UC]. I was attending NYU graph theory seminars and helped Clyde Kruskal proofread [UC]. This important paper proposed two models of parallel computation. The primary model, which Jack called the ultracomputer, was a message passing machine in which each processor-with-memory was connected to 4 others – the key connection being the perfect shuffle. In addition, he proposed an idealized model, the paracomputer, in which a large central memory could be simultaneously accessed in a single cycle by multiple independent processors. Although fan-in limitations would prevent this latter model from being realized, it has served as a standard against which others are compared. Soon Malcolm Harrison and Larry Rudolph become interested in Jack's paper and the ultracomputer group was unofficially born.

As luck would have it, the next year 1979-1980 was my sabbatical year from CUNY so I went to NYU to continue studying graph theory. However, the lure of an embryonic project on a fascinating subject was too strong. I switched to parallel processing. All the work from that early period concerned the ultracomputer model. Lambert Merteens, visiting Courant from the Mathematics Centrum, presented a stepwise refinement derivation of parallel bitonic sort [UCN1] and showed that large ultracomputers could not be built from small ones [UCN2]. Jack proposed one programming language [UCN3] and I wrote a simulator for another [UCN10] and [UCN14], which Clyde and I subsequently extended [UCN15]. We also wrote [UCN11], which extended the applicability of algorithms in [UC] and showed that permutations are hard to parallelize. The latter result is to appear as [PERM] and Clyde eventually parlayed this study into his thesis [CT]). Clyde and Larry [UCN6] developed an idea of Malcolm Harrison and Mal Kalos – multidimensional ultracomputers. Charles Peskin and Olof Widlund investigated scientific programming on the ultracomputer [UCN18], [UCN19], & [UCN20]. There was also a (short lived) flurry of activity in constructing “nearly planar” shuffle-exchange interconnections [UCN4], [UCN8], & [UCN9].

At this point the project was strongly influenced by an event that occurred outside NYU. Burroughs Corporation proposed an MIMD omega-network-based shared memory machine for the NASA Ames “digital wind tunnel” (NASF). We read their phase one proposal and traveled to Paoli PA to exchange ideas. Although we had hoped to have our work influence this potentially important commercial machine, in retrospect I believe that they affected us more than we affected them.

Jack, realizing that such a network-based machine could provide a realizable approximation to the paracomputer, wrote [UCN5], which provided a theoretical explanation of why the Burroughs machine attained good average case performance for data permutations and Clyde and I resurrected an ultracomputer permutation algorithm with good average case performance [UCN7] that we had discarded while working on [UCN11] (since its worst case performance is poor). The Burroughs design stimulated our interest in paracomputers both as theoretical models and as ideal machines toward which pragmatic machine designs should aim; for example, Clyde adapted the ultracomputer algorithms of [UCN11] to paracomputers [UCN26] (another part of [CT]).

About a dozen years earlier Ralph Grishman had written a software simulator of a machine quite similar to the paracomputer. This simulator, dubbed SOAPSUDS, was used for some early parallel processing studies at NYU [BU1] & [BU2] around 1970. After extensive conversations with Ralph (who subsequently became our first hardware designer), I wrote a paracomputer simulator called WASHCLOTH [UCN12] & [UCN21], which

semantics of the newly implemented operation were (trivially) different from those of the previous one, I eventually agreed with Mal and Clyde that we should switch to the (more accurate) name, "fetch-and-add".

The new year 1981-1982 brings us Pat Teller, Jim Lipkis, George Holober, and Jim Wilson as well as a strong desire to construct a prototype (NYU) Ultracomputer.

has since been heavily used for scientific programming [UCN22], [UCN23], [UCN24], [UCN27], [UCN30], [UCN31] (David Korn, on leave from Bell Labs 1980-81 did systems work on WASHCLOTH as well as scientific programming; Gabi Leshem was a graduate student working on the project; and Norman Rushfield helped us parallelize a large NASA "weather code"). I believe that the continual reliance on realistic scientific applications to test our theories has been a significant strength of the project.

The 1970 studies used a single "NEWVAL" primitive to coordinate the processors. Although NEWVAL was supported by SOAPSUDS, no one then had any idea how to implement it in a nonserial manner. We also used this primitive (but called it replace-add) in WASHCLOTH and hoped to find an efficient implementation on some paracomputer approximation like the Burroughs machine. Larry found a way to do this by enhancing the switches in the Burroughs design. Moreover, Larry, Boris Lubachevsky, and I found fully parallel algorithms for queue insertion and deletion. These last two discoveries, plus additional fully parallel software primitives, were the subject of [UCN16] and resulted in our current emphasis on network-based designs since we came to realize that most other architectures do not permit such fully parallel algorithms (for example see [UCN17] which has appeared as [TC]). An abbreviated version of [UCN16] appeared as [PP], essentially the entire work is to appear as [TOPLAS], and Larry parlayed it into [LT]. After writing a program to check several algorithms in [UCN16], Boris became interested in automatic verification of parallel programs, eventually producing [UCN33].

This past year 1980-81 began with the arrival of Marc Snir and saw Kevin McAuliffe become actively involved. As the year closed the entire group (expanded by the mid-year arrival of Shaula Yemini) realized that we should seriously consider building a prototype. During that year Marc analyzed general connection networks [UCN29] and produced a modified WASHCLOTH that introduced delays based on his queuing theory approximation of our proposed network [UCN28]; Boris surveyed the Russian literature in parallel processing [UCN35]; Jack and I summarized the group's network-based work [UCN25] which appeared as [C]; and I wrote a toy parallel operating system MOP [UCN13] on top of WASHCLOTH to test several software primitives in [UCN16]. By then we had a definite macroscopic machine design as well as a detailed design for selected components (e.g. the switches). We presented this design in [UCN32] (a nearly final version of which appeared as [SD]) and named it the NYU Ultracomputer on the grounds that what else could you call a machine designed by the NYU Ultracomputer project (definition "Mathematics"— the subject matter studied by mathematicians). Finally, Clyde and I found a new parallel implementation of replace-add [UCN34]. This implementation also relied on enhanced switches; in fact it was quite similar to the [UCN16] implementation (but generalized to other important operations). Since the

Bibliography

[BU1]

E. Draughon, Ralph Grishman, J. T. Schwartz, and A. Stein, "Programming Considerations for Parallel Computers", Courant Institute Report IMM 362, November 1967.

[BU2]

E. Draughon, J. T. Schwartz, and A. Stein, "Individual and Multi-Processing Performance Characteristics of Programs" on Large Parallel Computers", Courant Institute Report IMM 380, April 1970.

[C] Allan Gottlieb J. T. Schwartz, "Networks and Algorithms for Very Large Scale Parallel Computation", *Computer*. 15 pp. 27-36, January 82.

[CT]

Clyde P. Kruskal, "Upper and Lower Bounds on the Performance of Parallel Algorithms", Thesis, NYU, October 1981.

[LT]

Larry Rudolph, "Software Structures for Large-Scale Parallelism", Thesis, NYU, to appear.

[PERM]

Allan Gottlieb and Clyde P. Kruskal, "Complexity Results for Permuting Data on Parallel Processors", submitted for publication.

[PP] Allan Gottlieb, B. D. Lubachevsky, and Larry Rudolph, "Coordinating Large Numbers of Processors", *Proc. 1981 Internat. Conf. on Parallel Processing*, Bellaire MI, pp. 341-349, August 1981.

[SD]

Allan Gottlieb, Ralph Grishman, Clyde P. Kruskal, Kevin McAuliffe, Larry Rudolph, and Marc Snir, "The NYU Ultracomputer - A General Purpose Parallel Processor", *Proc. Signal Processing IV at SPIE Annual Conf.*, San Diego, August 1981.

[TC]

Allan Gottlieb, "Comments on 'Concurrent Search and Insertion in AVL Trees'", *IEEE Trans. on Computers* 30 p. 812, October 1981.

[TOPLAS]

Allan Gottlieb, B. D. Lubachevsky, and Larry Rudolph, "Basic Techniques for the Efficient Coordination of Very Large Numbers" of Cooperating Sequential Processors", to appear in *ACM TOPLAS*.

[UC]

J. T. Schwartz, "Ultracomputers", 2, #4, pp. 484-521, October 1980.

[UCN1]

Lambert Meertens, "Bitonic Sort on Ultracomputers", March 1979.

- [UCN2]
Lambert Meertens, "Recurrent Ultracomputers are not $\log N$ -Fast",
March 1979.
- [UCN3]
Jacob T. Schwartz, "Preliminary Thoughts on Ultracomputer
Programming Style", May 1979.
- [UCN4]
Jacob T. Schwartz, "A Remark on Nearly Planar Embeddings of Small
Ultracomputers", December 1979.
- [UCN5]
Jacob T. Schwartz, "The Burroughs FMP Machine", January 1980.
- [UCN6]
Clyde P. Kruskal and Larry Rudolph, "Observations Concerning
Multidimensional Ultracomputers", January 1980.
- [UCN7]
Allan Gottlieb and Clyde P. Kruskal, "A Data-Motion Algorithm",
January 1980.
- [UCN8]
Larry Rudolph, "A Remark on the Planarity of the Shuffle-Exchange
Network" of Sizes 16 and 32", February 1980.
- [UCN9]
Allan Gottlieb, "Another Remark on the Planarity of Shuffle-Exchange
Networks" of Sizes 16 and 32", May 1980.
- [UCN10]
Allan Gottlieb, "PLUS - A PL/I Based Ultracomputer Simulator, I",
September 1980.
- [UCN11]
Allan Gottlieb and Clyde P. Kruskal, "Supersaturated Ultracomputer
Algorithms", September 1980.
- [UCN12]
Allan Gottlieb, "WASHCLOTH - The Logical Successor to Soapsuds",
October 1980.
- [UCN13]
Allan Gottlieb, "MOP - A (Minimal) Multiprocessor Operating System
Extending WASHCLOTH", November 1980.
- [UCN14]
Allan Gottlieb, "PLUS - A PL/I Based Ultracomputer Simulator, II",
November 1980.
- [UCN15]
Allan Gottlieb and Clyde P. Kruskal, MULT - A Multitasking

Ultracomputer Language with Timing, I & II", December 1980.

[UCN16]

Allan Gottlieb, B. L. Lubachevsky, and Larry Rudolph, "Basic Techniques for the Efficient Coordination of" Very Large Numbers of Cooperating Sequential Processors", December 1980.

[UCN17]

Allan Gottlieb, "Comments on 'Concurrent Search and Insertion in AVL Trees'", December 1980.

[UCN18]

Charles S. Peskin and Olof B. Widlund, "Remarks on Efficient Numerical Algorithms for Ultracomputers", January 1981.

[UCN19]

Charles S. Peskin, "Ultracomputer Implementation of Odd-Even Reduction", January 1981.

[UCN20]

Charles S. Peskin, "A Comparison of Ultracomputer Architecture and Lattice" Architecture for Problems on Lattices", January 1981.

[UCN21]

Allan Gottlieb, "WASHCLOTH #81", January 1981.

[UCN22]

Norman Rushfield, "Atmospheric Computations on Highly Parallel MIMD Computers", February 1981.

[UCN23]

David Korn, "Converting Scientific Codes to Run Under the WASHCLOTH Simulator", March 1981.

[UCN24]

David Korn, "Timing Analysis for Codes Run Under the WASHCLOTH Simulator", March 1981.

[UCN25]

Allan Gottlieb and Jacob T. Schwartz, "Networks and Algorithms for Very Large Scale Parallel Computation", March 1981.

[UCN26]

Clyde P. Kruskal, "Supersaturated Paracomputer Algorithms", May 1981.

[UCN27]

Malvin Kalos, Gabi Leshem, and B.D. Lubachevsky, "Molecular Simulations of Equilibrium Properties", May 1981.

[UCN28]

Marc Snir, "NETSIM Network Simulator for the Ultracomputer", May 1981.

- [UCN29]
Marc Snir, "Lower Bounds on VLSI Implementations of Communication Networks", May 1981.
- [UCN30]
Malvin Kalos, "Scientific Calculations on the Ultracomputer", June 1981.
- [UCN31]
David Korn, "Timing Simulations for Elliptic PDE's Run Under WASHCLOTH", June 1981.
- [UCN32]
Allan Gottlieb, Ralph Grishman, Clyde P. Kruskal, Kevin P. McAuliffe, Larry Rudolph, and Marc Snir, "The NYU Ultracomputer - A General-Purpose Parallel Processor", July 1981.
- [UCN33]
B. D. Lubachevsky, "Verification of Several Parallel Coordination Primitives Based on" Descriptions of Their Reachability Sets", July 1981.
- [UCN34]
Allan Gottlieb and Clyde P. Kruskal, "Coordinating Parallel Processors: A Partial Unification", September 1981.
- [UCN35]
B. D. Lubachevsky, "Review of Soviet Publications on Parallel Data Processing (1978-80)", September 1981.
- [UCN36]
Allan Gottlieb, "A Historical Guide to the Ultracomputer Literature", October 1981.

